




## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<b>(51) International Patent Classification <sup>6</sup>:</b> <b>H04L 12/407, 12/413, G06F 13/376</b>	<b>A1</b>	<b>(11) International Publication Number:</b> <b>WO 99/43130</b> <b>(43) International Publication Date:</b> 26 August 1999 (26.08.99)
<b>(21) International Application Number:</b> PCT/US99/03070 <b>(22) International Filing Date:</b> 12 February 1999 (12.02.99) <b>(30) Priority Data:</b> 09/026,884                      19 February 1998 (19.02.98)                      US <b>(71) Applicant:</b> EPIGRAM, INC. [US/US]; 870 West Maude Avenue, Sunnyvale, CA 94086 (US). <b>(72) Inventors:</b> HOLLOWAY, John, T.; 7 Oakhill Drive, Woodside, CA 94062 (US). TRACHEWSKY, Jason; Apartment C, 854 Fremont Street, Menlo Park, CA 94025 (US). PTASINSKI, Henry; Apartment 2, 1638 Latham Street, Mountain View, Ca 94041 (US). <b>(74) Agents:</b> EDWARDS, Gary, J. et al.; Skjerven, Morrill, MacPherson, Franklin & Friel LLP, Suite 700, 25 Metro Drive, San Jose, CA 95110 (US).		<b>(81) Designated States:</b> AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).  <b>Published</b> <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

**(54) Title:** A PACKET-SWITCHED MULTIPLE-ACCESS NETWORK SYSTEM WITH DISTRIBUTED FAIR PRIORITY QUEUING

**(57) Abstract**

A packet-switched multiple-access network system with a distributed fair priority queuing media access control protocol that provides multiple levels of priority of access and fair collision resolution with improved performance is disclosed. In one embodiment, the system provides high-speed transport of multimedia information on a shared channel. Further, in one embodiment, MAC level side-band signaling that is useful to other levels of the network protocol (e.g., the physical layer) is also provided.

**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

## A PACKET-SWITCHED MULTIPLE-ACCESS NETWORK SYSTEM WITH DISTRIBUTED FAIR PRIORITY QUEUING

### FIELD OF THE INVENTION

5           The present invention relates to packet-switched communication systems, and in particular, to an apparatus and a method for a packet-switched multiple-access network system with distributed fair priority queuing.

### BACKGROUND

10           Packet-switched communication systems (e.g., packet-switched networks) are useful for the transport of data over shared communication channels on a variety of physical media such as twisted-pair, coaxial cable, power lines, and wireless. The data communicated may include multimedia information such as voice data and video data.

15           One such system that is in wide-spread commercial use is standard ethernet (e.g., ethernet is commonly used in LANs (Local Area Networks)). Systems such as standard ethernet use a multiple-access technique to coordinate access among several stations contending for use of a shared channel. In particular, standard ethernet is based on 1-persistent Carrier Sense Multiple Access with Collision  
20 Detect (CSMA/CD) with a Collision Resolution Algorithm (CRA) referred to as Binary Exponential Backoff (BEB).

          In a CSMA/CD network, a shared communication channel is shared by network terminals or stations (e.g., personal computers and printers).  
Transmissions on the channel are segmented into variable length packets. Only  
25 one station is granted access to place (i.e., transmit) a packet on the channel at any given time, and the presence of a packet can be detected by all the stations using a carrier sensing device. All stations obey a distributed access protocol that includes the following stages: (1) a station wishing to transmit a packet monitors the channel to detect the absence of carrier; (2) when the channel is idle each

contending station commences transmission; and (3) if a transmitting station detects energy from another transmitting station during its use of the channel (i.e., collision detection), then the station abandons its transmission and activates a Collision Resolution Algorithm (CRA) to resolve the access ordering among the  
5 contending stations.

In standard ethernet, the CRA is BEB. In BEB, a count of the number of collisions (N) during attempts to transmit a given packet is maintained by each station, a random number K is chosen from the interval 0 to  $2^N-1$ , and the station waits for slot K following the end of the current transmission before attempting  
10 transmission. If some other station commences transmission before slot K the current station defers until the end of that transmission and restarts its CRA.

However, the standard ethernet protocol lacks robust performance and is inefficient. For example, the BEB approach can cause an order of magnitude increase in channel access latency for some stations under modest offered load  
15 (also known as the packet starvation effect in CSMA/CD LANs).

#### SUMMARY

Accordingly, the present invention provides a packet-switched multiple-access network system with a Distributed Fair Priority Queuing (DFPQ) MAC (Media Access Control) protocol that provides improved performance and  
20 efficiency. In one embodiment, the present invention provides a fair collision resolution MAC protocol with multiple priority levels of access. In this embodiment, a stack or tree algorithm is used instead of the BEB algorithm of standard ethernet.

In particular, the time period following the end of a previous transmission  
25 is divided into slots, which are sized such that a signal can traverse both directions of the longest path in the shared channel with some margin during one slot. In particular, when used for contention, the slots are numbered by priority level with the highest priority level first. A first station selects the highest priority traffic waiting in the first station's transmit queue and contends during the slot assigned to

that priority. Higher priority traffic from a second station will have commenced transmission before the slot corresponding to the first station's traffic and will therefore have absolute priority of access. When no higher priority traffic is waiting, the first station will have priority over all lower priority traffic waiting in  
5 other stations as the first station will commence transmission and defer the other lower priority stations.

If multiple stations have traffic waiting at the same priority level, then the multiple stations will attempt transmission simultaneously and collide. The collision is detected by transmitting stations and signaled to all stations. The  
10 collision can be detected, for example, by the fact that the collision transmission ends up being shorter in length than a legitimate data packet. Upon detection of a collision, all stations enter a contention resolution phase, and specifically, a stack-based collision resolution cycle in accordance with one embodiment of the present invention. More specifically, a fair collision resolution MAC protocol is provided  
15 in which the first S slots following a collision are used as signal slots. For example, three signal slots can be provided. Upon a collision, each active station rolls an S-sided die to vote for one of the S signal slots (i.e., randomly selects a number between 0 to S) and transmits a signal in the chosen slot. Each station maintains a stack counter for each priority level. The priority level of the active  
20 collision resolution phase is the priority level of the collision that commenced the active collision resolution phase. All active stations monitor the signal slots, and for each signal preceding the slot chosen by a given station, the station increments its stack counter. Accordingly, each station that started the phase at a non-zero stack level increments its stack counter for each signal detected. After a successful  
25 transmission without collision, each active station decrements its stack counter for that priority level. In this manner, the stations order themselves at different stack levels based on the random vote of signal slots in each phase. In this embodiment, if another collision occurs in the subsequent contention slot, then the collision resolution protocol is re-invoked.

In this embodiment, each priority level has a separate stack counter, and a collision resolution sequence at a given priority level can be interrupted by transmissions at a higher priority level at any time. Also, the signal slots are only present immediately after a collision and do not create overhead after successful  
5 transmissions. In addition, unlike conventional approaches, slots are either used for priority ordering, for sensing a collision, or sending a signal, and the different uses are optimally selected to minimize overhead on the channel.

Further, in one embodiment, if the channel has been idle for a time longer than the series of slots assigned to defined priority levels, then all stations are free  
10 to contend. In this case, transmission order is determined on a first-come, first-served basis (i.e., no station has absolute priority).

Simulations have been performed examining the optimal number of signal slots (e.g., whether to use two signal slots, three signal slots, four signal slots, etc.). Under a set of assumptions regarding relative sizes of packets and slots, three  
15 signal slots was shown to be optimal, although the system performed nearly as well with two or four signal slots.

Further, in one embodiment, stations are assigned priority levels according to bandwidth class such that higher baud rate transmissions are given preferential priority to the segment. As a result, the aggregate capacity of the segment is  
20 optimized towards the higher baud rate.

In another embodiment, stations alternate between priority levels with probability according to bandwidth class such that the aggregate capacity is divided in any proportion between different baud rate transmissions. As a result, the aggregate capacity is optimized and access is fairly apportioned between stations of  
25 different bandwidth class.

In another embodiment, instead of choosing a random slot at each contention resolution phase, a deterministic choice can be made, and furthermore, a deterministic contention resolution sequence of arbitrary depth can be forced by short transmissions mimicking collisions. In particular, the sequence of

deterministic choices finalizing in a resolution represent a label or side-band data (e.g., side-band signaling) conveyed entirely within the MAC layer. In this embodiment, the label is used to identify the source station in multicast transmissions to assist the physical (PHY) layer in choosing demodulation parameters appropriate to the multicast source. Further, an additional use of the MAC labeling mechanism is to allow stations without the ability to detect collisions (e.g., which may occur during initialization on channels that present strong echoes), to gain access to the channel for echo probing. By sending a sufficiently long and unique label using the MAC contention resolution signaling described above, an initializing station can ensure an arbitrarily low probability of collision with another station, even though the initializing station lacks the ability to detect collisions.

In yet another embodiment, a HOLDOFF signal is appended to the end of a successful packet transmission, which signals all the distributed MAC entities to delay contention for a block of slots. For example, the HOLDOFF signal is useful for accommodating exceptional round-trip latencies between a central office station remote from LAN stations.

Other aspects and advantages of the present invention will become apparent from the following detailed description and accompanying drawings.

## BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a shared-medium broadcast network with attached stations in accordance with one embodiment of the present invention.

FIG. 2 is a functional diagram of the network interface of FIG. 1 in accordance with one embodiment of the present invention.

FIG. 3 is a functional diagram of the multiple access controller of FIG. 2 in accordance with one embodiment of the present invention.

FIG. 4 is a timing diagram showing packet transmissions, slots, and collisions on a shared channel of the shared-medium of FIG. 1 in accordance with one embodiment of the present invention.

FIG. 5 is a diagram of a tree that represents a sequence of choices of signal slots selected during a contention cycle in accordance with one embodiment of the present invention.

FIG. 6 is a plot of maximum and minimum latency versus offered load  
5 comparing the performance of a Distributed Fair Priority Queuing (DFPQ) Media Access Control (MAC) protocol in accordance with one embodiment of the present invention with standard ethernet. FIG. 7 is a timing diagram showing packet transmissions and slots of the shared-medium of FIG. 1 in accordance with another embodiment of the present invention.

10 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

In one embodiment, the present invention provides a packet-switched multiple-access network system with a Distributed Fair Priority Queuing (DFPQ) MAC (Media Access Control) protocol that provides improved performance and also provides fair collision resolution with multiple priority levels of access. For  
15 example, a network that offers guarantees of Quality of Service (QoS) to certain traffic streams such that controlled streams access the network at a higher priority than uncontrolled traffic would significantly benefit from this embodiment of the present invention.

FIG. 1 is a block diagram of a shared-medium broadcast network 101 with  
20 attached stations (102, 104, 105, and 106) in accordance with one embodiment of the present invention. In particular, the physical medium of shared-medium 101 is either: (1) a network of connected twisted-pair wiring segments such as found in premises telephone wiring; (2) a network of power wiring such as found within a residence; (3) a network of connected coaxial cable segments such as found in  
25 premises CATV (cable television) wiring; (4) a wireless medium represented by infrared, RF (Radio Frequency) or microwave broadcast; or (5) any combination of the above.

Station 102 includes a network interface 103 that attaches the host logic of station 102 to shared-medium 101. In particular, network interface 103 transmits



and receives packets of information. In this embodiment, all stations (i.e., stations 102, 104, 105, and 106) include a network interface such as network interface 103. The access control for shared-medium 101 is accomplished through a DFPQ (Distributed Fair Priority Queuing) MAC (Media Access Control) protocol  
5 operating in each network interface in each station. Thus, there is no central controller.

The signals on shared-medium 101 may be baseband (using spectrum extending up from DC) or passband (using a spectral range centered away from DC). The signals may use a range of modulation and coding techniques to  
10 communicate data within a packet such as (1) pulse amplitude modulation; (2) quadrature amplitude modulation and carrierless amplitude phase modulation; (3) multi-tone modulation; or (4) frequency or phase modulation.

In particular, the signals using these modulation and coding techniques can be assigned to logical channels on a given physical medium using multiplexing  
15 techniques such as (1) baseband, utilizing frequencies down to DC; (2) frequency division multiplexed, wherein each channel uses a distinct passband of frequencies on the physical medium; (3) code division multiplexed, wherein each channel is associated with a set of pseudo-random code sequences; or (4) time division multiplexed, wherein each channel occupies a specific time interval in a repeating  
20 time frame. For example, if employing passband modulation, there may be a series of logical channels Frequency Division Multiplexed (FDM) onto shared-medium 101. In this example, each logical channel may have a unique instance of a multiple access control entity, which is discussed below with respect to FIG. 2. The logical channel used for the specific signals comprising the distributed fair  
25 priority queuing MAC protocol may be shared with the channel used for payload data of packet transmissions, or alternatively the MAC specific signals can occupy a separate logical channel coordinated with data transmissions in other logical channels. In a preferred embodiment, a single channel passband Quadrature Amplitude Modulation (QAM) technique is used.

FIG. 2 is a functional diagram of network interface 103 in accordance with one embodiment of the present invention. In particular, attachment to shared-medium 101 is made via a HYBRID 209 which separates the receive signal (into a demodulator 203) from the transmit signal (from a modulator 205). A demodulator 203 includes a signal chain that includes gain control, matched filter, resampler, feed-forward equalizer, decision feedback equalizer, demodulator, timing recovery for carrier and baud, sampler, maximum likelihood sequence estimator, error decoder, and framing logic. Demodulator 203 converts the continuous time signals from HYBRID 209 into packet data which is stored in an RX queue 207 for use by the host station logic of station 102.

A carrier & signal detect 202 computes the power level of the received signal and compares this level to a threshold level producing a binary indication of a presence or absence of signal. The output of carrier & signal detect 202 is connected to a multiple access controller 206.

A modulator 205 generally performs the inverse function of demodulator 203. In particular, modulator 205 converts packet data from a TX queue 208 into a continuous time signal that can be applied to shared-medium 101. Control of the modulator 205 transmission is made by multiple access controller 206. The transmitted output of modulator 205 is compared with the received signal in a collision detect 204. Any detected difference indicates that some other station is transmitting on shared-medium 101. In a preferred embodiment, the detection operation involves computation of an echo replica of the transmitted signal using a standard FIR filter calibrated with coefficients measured during non-colliding transmissions. The power level of the detected difference signal is compared to a threshold level and generates a collision detect signal which is connected to multiple access controller 206.

Multiple access controller 206 also receives input from TX queue 208 regarding the priority such as priority signals 201 and desired (MAC) label (if any) for the packet ready to be transmitted. Multiple access controller 206 also outputs

the detected priority and label of received signals to RX queue 207 along with the packet data from demodulator 203.

Simulations have been performed examining the optimal number of signal slots (e.g., whether to use two signal slots, three signal slots, four signal slots, etc.).

- 5 Under a set of assumptions regarding relative sizes of packets and slots, three signal slots was shown to be optimal, although the system performed nearly as well with two or four signal slots.

FIG. 3 is a functional diagram of multiple access controller 206 in accordance with one embodiment of the present invention. In particular, inputs to  
10 CONTROL LOGIC 301 include a carrier sense signal from carrier & signal detect 202 as well as a collision detect signal from collision detect 204. Also, the priority level and desired label (if any) of the packet in TX queue 208 that is to be transmitted are input to CONTROL LOGIC 301.

When the carrier sense signal becomes inactive, CONTROL LOGIC 301  
15 activates a SLOT TIMER 302, which identifies slots as contentions or signals. In particular, one function of CONTROL LOGIC 301 is to monitor for collisions and signals and to maintain a counter STACK 303 and a counter STACK TOP 304. The following pseudo-code illustrates an implementation of this functionality of CONTROL LOGIC 301 in accordance with one embodiment of the present  
20 invention:

```

After successful packet transmission,
    STACK[priority]--, STACKTOP[priority]--
ELSE collision THEN
25     IF STACK[priority]=0, THEN select S0, S1, or S2, send signal,
        IF selected S1, THEN STACK[priority] += S0
        IF selected S2, THEN STACK[priority] += S0 + S1
        ELSE STACK[priority] += S0 + S1 + S2 - 1
    STACKTOP[priority] += S0 + S1 + S2 - 1
30     IF new packet arrival, THEN STACK[priority] = STACKTOP[priority]+1.
  
```

CONTROL LOGIC 301 will only enable transmission (e.g., station active) when STACK counter 303 at the current priority level is equal to zero. If the given station has encountered a collision on previous attempts, it may be that STACK  
5 counter 303 has been increased above zero thereby making the given station defer to other contending stations. At the end of a successful transmission, all stations decrease their STACK counter 303 by one, which will make those stations at the top of the stack (counter value one) now active.

In one embodiment, when a station encounters a collision, the station  
10 randomly selects a value in the range 0...2 (i.e., assuming that there are three signal slots: S0, S1, and S2) and asserts a signal in the corresponding signal slot. This signal informs all the contending stations that there are one or more active stations choosing that slot. These signals allow the stations to resolve their ordering. In particular, those stations signaling in lowered numbered slots will contend before  
15 stations selecting higher numbered slots. The STACK counter 303 keeps track of this order. If a station is not active (i.e., STACK counter 303 is non-zero), then STACK counter 303 increases its stack level to accommodate previously active stations now entering the stack. Only the stations signaling in the lowest numbered signal slot will be active in the next round.

20 Counter STACK TOP 304 tracks the highest stack level for any station, at each priority level. This is used to set the stack level for stations becoming active with new packet data during a cycle of collision resolution. This improves the fairness of access, ensuring that late arriving stations do not gain preference over stations already contending. In an alternative embodiment, a "closed" entry policy  
25 is used, which inhibits new packet transmission until all the currently contending stations have finished transmission. In simulations, this alternative embodiment was found to produce non-optimal latency distributions.

In a preferred embodiment, the above pseudo-code is mapped into a hardware description language such as Verilog or VHDL, and automated synthesis  
30 tools (e.g., commercially available automated synthesis tools from Synopsys) are

used to compile the computational structure into an ASIC (Application Specific Integrated Circuit). Alternatively, the pseudo-code can be translated into a high level programming language such as "C" and executed on a standard embedded microcontroller that is part of network interface 103. Alternatively, the pseudo-code can be translated into firmware and executed on a custom microcontroller or hardware finite-state machine.

FIG. 4 is a timing diagram showing packet transmissions, slots, and collisions on a shared channel of shared-medium 101 in accordance with one embodiment of the present invention. A successful transmission is shown as PACKET 401. Immediately following the end of packet 401, a timer is started which divides time periods into a series of slots. The slots are ordered by priority with the highest priority first.

For example, in FIG. 4, eight (8) priority levels are shown, numbered 7...0. Priority level 6, slot 402, is the second slot following the end of PACKET 401. During slot 402 all stations with traffic ready at priority 6 attempt transmission. If only one such station is active, a successful packet transmission is made (e.g., PACKET 405).

However, if more than one station contends, then a collision is sensed (by collision detect 204 in each transmitting station). For example, the hybrid 209 separates the received signal from the transmitted signal produced by modulator 205, and a collision is detected by the presence of received energy from other transmitting stations. In another embodiment, the collision detect 204 augments the hybrid by subtracting an echo replica of the transmitted signal.. In FIG. 4, a collision on the shared channel is indicated by reference numeral 406. Immediately following the collision, a series of signal slots are used to signal between colliding stations to resolve the collision. These slots are shown as S0, S1 (identified by reference numeral 404), and S2. Immediately following the sequence of signal slots, contention slots restart at priority 7. Assuming no priority 7 station has become active, then a subset of the stations contending in slot 402 retry transmission in slot 408. If the number of contending stations has been

reduced to one, a successful packet is sent (e.g., PACKET 407), otherwise another collision occurs as indicated by reference numeral 403, and the collision resolution protocol is re-invoked.

Further, the above described (access) priority mechanism can be used to  
5 segregate traffic on a given segment (e.g., shared channel) that uses disparate modulation rates, which is particularly advantageous in a scaleable network system to allow mixed transmission at different baud rates. For example, a low cost interface may use transmissions at a rate of 1 Mbaud/sec (Megabaud per second), whereas a high performance interface may use transmissions at a rate of 10  
10 Mbaud/sec. In this example, there is the possibility that traffic offered at the lower baud rate may occupy ten times as much time transmitting on the segment as traffic offered at the higher baud rate, and therefore, the aggregate bandwidth capacity of the segment is detrimentally reduced. Accordingly, in one embodiment, stations can be assigned (access) priority levels according to bandwidth class such that  
15 higher baud rate transmissions are given preferential priority to the segment. As a result, the aggregate capacity of the segment is optimized towards the higher baud rate. In another embodiment, stations choose between (access) priority levels with a statistical distribution designed to optimize the utilization of the aggregate capacity while ensuring that transmissions at a lower bandwidth class are not  
20 unfairly excluded from access to the channel.

In another embodiment, a HOLDOFF signal is inserted at the end of each successful transmission interval, at a time just prior to the contention slots (e.g., at a time identified by reference numeral 409 of FIG. 4). This signal informs the distributed multiple access controller 206 entities to insert a block of slots (e.g., 16  
25 slots) before starting the countdown of contention slots. This time interval is then made available to a remote station (e.g., located at a central office at the other end of a wide-area subscriber loop) to commence transmission without superposing the requirement that the remote station be able to participate in the contention protocol (i.e., execute the DPFQ MAC protocol). This approach also allows multiple access

controller 206 to synchronize contention cycles to an overall frame interval by detecting the proximity to a frame boundary and commencing a transmission which is artificially terminated to coincide with the origin of the next frame. Assertion of the HOLDOFF signal at the end of this special transmission can then  
5 allow the remote station to commence transmissions at a known time. For example, this network frame clock is useful in certain applications for sending synchronized data.

FIG. 5 is a diagram of a tree 500 that represents a sequence of choices of signal slots selected during a contention cycle in accordance with one embodiment  
10 of the present invention. In particular, the sequence of choices of signal slots chosen during a contention cycle can be thought of as an address or (MAC) label. At the start of the cycle, shown as a node 501 at the top of tree 500, a station chooses one of three paths, shown as S0 (502), S1, or S2. The choice leading to the next node of tree 500 forms a digit of the label 0xx (503). As shown in FIG. 5,  
15 the next two selected choices are S1 leading to a label of 01x (504) and S1 leading to a label of 011 (505).

In one embodiment, these choices are selected randomly to ensure that the expected depth of the tree needed to resolve to a single station is optimally minimum. In this embodiment, the derived label merely reflects the random  
20 choices selected and conveys no useful information.

However, if a station makes non-random selections at each level of the contention tree, then when the station finally resolves successfully and transmits its packet, the station has also conveyed information in the form of a (MAC) label to the receiving station(s). In a preferred embodiment, the label is used to provide  
25 side-band signaling in four ways:

1) The label is used by multicasting stations to identify the source of a multicast and assist in the PHY layer. In particular, each multicasting station has been assigned a unique label, and the receiving stations use the label information to select the demodulation parameters corresponding to the identified source;

2) The label is used to select a unique sub-priority within a given priority level;

3) The label is used to force a low probability of collision for stations that cannot yet detect collisions, but need access to shared-medium 101 to transmit initialization signals to calibrate their collision detect blocks 204. In particular, the MAC labeling mechanism is used to allow stations without the ability to detect collisions (e.g., which may occur during initialization on channels that present strong echoes) to gain access to the channel for echo probing. By sending a sufficiently long and unique label using the MAC contention resolution signaling described above, an initializing station can ensure an arbitrarily low probability of collision with another station, even though the initializing station lacks the ability to detect collisions; and

4) The label is used to identify certain packets for special treatment by layers or sub-layers of the protocol stack other than the PHY layer.

FIG. 6 is a plot of maximum and minimum latency versus offered load comparing the DFPQ MAC protocol in accordance with one embodiment of the present invention with standard ethernet. In particular, FIG. 6 shows a simulation of ten stations with uniformly distributed traffic load in which the offered load was increased in steps, and the minimum and maximum of average latency across all ten stations is recorded for each run. As the offered load approaches 100%, ETHERNET MAX LATENCY shows the unfair access statistics known as the packet starvation effect, as discussed above. In particular, some stations see continuously increasing delay (601) while other stations "capture" the shared-medium and see minimum latency (603). In contrast, DFPQ MIN/MAX LATENCY shows a tight bound on latency (602) and is stable even under offered loads greater than 100%.

FIG. 7 is a timing diagram showing packet transmissions and slots of shared-medium 101 in accordance with another embodiment of the present invention. In particular, the addition of a deterministic label indicator to the slot structure of the present invention allows a listening station to differentiate between



a frame transmitted with a randomly-selected sequence of signal values ( $S$  choices, as discussed above with respect to FIG. 5) and a frame transmitted with a deterministic choice of signal values. Hence, a station that desires to transmit a deterministic CRA sequence also signals in a slot 701 (referred to as a "DI" slot) immediately following slot  $S(n)$ , where  $n$  is the number of signal slots, and before the priority  $p$ , where  $p$  is the number of transmission priorities, transmission slot. The station only signals in DI slot 701 when it has transmitted the entire CRA sequence corresponding to the desired label and when it has chosen the first occupied signal slot (i.e., no other station chose slot  $S(k-1)$  if this station chose slot  $S(k)$ ). There are three possible outcomes of this situation: (1) the station which has transmitted the deterministic label successfully transmits in its slot; (2) the station is preempted by a higher priority transmission; and (3) the station collides with another station of the same priority. In cases (2) and (3), signaling in the DI slot has no effect. However, in case (1), signaling in the DI slot serves to inform the listening stations that the transmission is one labeled by a deterministic CRA sequence. Thus, a listening station can unambiguously set the correct PHY parameters for the source.

Accordingly, in one embodiment, the present invention provides an efficient absolute access priority protocol that can be enforced per packet per station. The latency for higher priority traffic is almost entirely insulated from the offered load at lower priorities. In particular, multimedia traffic can be assured bounded access latencies proportional to the effective number of active stations at the same or higher priority level, which is advantageous for providing guarantees of QoS. Also, the priority mechanism adds little overhead to the standard ethernet MAC layer. Further, the maximum latency seen by any station is fairly distributed, has low variation, and is bounded close to the theoretical minimum for a perfectly fair queue. In particular, the packet starvation phenomenon of standard ethernet is eliminated. Further, the MAC protocol can be used to balance usage of the medium by different bandwidth classes of transmission thereby maintaining a

higher aggregate capacity of the medium. Moreover, in one embodiment, the present invention provides additional signaling that is useful for other layers of the network protocol stack (e.g., the PHY layer).

Although particular embodiments of the present invention have been shown  
5 and described, it will be obvious to those skilled in the art that changes and modifications may be made without departing from the present invention in its broader aspects, and therefore, the appended claims are to encompass within their scope all such changes and modifications that fall within the true scope of the present invention.

CLAIMS

What is claimed is:

1. A packet-switched multiple-access network system comprising a  
5 shared communication channel, a first station connected to the shared  
communication channel, and a second station connected to the shared  
communication channel, wherein the first station and the second station each  
comprise:

10 a network interface, the network interface transmitting packets to  
the channel and receiving packets from the channel; and

a distributed fair priority queuing MAC (Media Access Control)  
protocol executed by the first station and by the second station, the  
distributed fair priority queuing MAC protocol resolving packet collisions  
between the first station and the second station on a per packet per access  
15 priority basis.

2. The system as in Claim 1 wherein the network interface further  
comprises:

20 a multiple access controller, the multiple access controller executing  
the distributed fair priority queuing MAC protocol.

3. The system as in Claim 2 wherein the network interface further  
comprises:

25 a carrier and signal detect that communicates with the multiple  
access controller, the carrier and signal detect detecting the presence or  
absence of transmissions on the channel.

4. The system as in Claim 3 wherein the network interface further comprises:

a demodulator, the demodulator converting received transmissions into packet data.

5

5. The system as in Claim 4 wherein the network interface further comprises:

a modulator, the modulator converting packet data into signals for transmission, wherein the transmission of the station is controlled by the multiple access controller using carrier sense.

10

6. The system as in Claim 5 wherein the network interface further comprises:

a collision detect that communicates with the multiple access controller, the collision detect detecting the presence of a simultaneous transmission (a collision).

15

7. The system as in Claim 6 wherein time intervals between packet transmissions over the channel are divided into a plurality of contention slots, and the contention slots are used for packet transmission contention opportunities that are ordered according to access priority.

20

8. The system as in Claim 7 wherein the signal slots are used for signaling opportunities for the first station and the second station to share information used for collision resolution by the distributed fair priority queuing MAC protocol.

25

9. The system as in Claim 8 wherein the signal slots are used for signaling opportunities for the first station and the second station to share information used for collision resolution implemented using a tree or stack-based collision resolution technique.

10. The systems as in Claim 8 wherein each station contending for access in a given collision resolution cycle pseudo-randomly selects among the possible signal slots and signals that vote by transmitting a signal in the selected signal slot.

11. The system as in Claim 8 wherein the network interface further comprises:

a first counter per priority that indicates the stack level above current active stations with which the first station is waiting to resolve its collision, and

control logic executed by the first station and the second station for incrementing and decrementing the first counter based on one or more of the following: (a) a carrier sense signal; (b) a collision detect signal; or (c) voting signals.

12. The system as in Claim 11 wherein the network interface further comprises a second counter that indicates the maximum stack level per priority of the first station and the second station, wherein the second counter is used to initialize the first counter when the first station has prepared a new packet for transmission.

13. The system as in Claim 8 wherein the pattern of votes used in a collision resolution cycle conveys side-band signaling information from the first station to the second station, and a deterministic label indicator slot is used to indicate that the pattern of votes is the result of the first station intending to convey  
5 side-band signaling information.

14. The system as in Claim 8 wherein the pattern of votes used in a collision resolution cycle conveys side-band signaling information from the first station to the second station, where the pattern of votes is used to specify a  
10 subordinate level of priority within the priority level associated with the collision resolution cycle.

15. The system as in Claim 8 wherein the pattern of votes used in a collision resolution cycle is chosen such that the probability of any second station  
15 further colliding with the first station is reduced.

16. The system as in Claim 8 wherein the collision resolution cycle comprises two, three, or four signal slots.

20 17. The system as in Claim 8 wherein the plurality of contention slots, which are ordered according to access priority, are further used for segregating transmissions at different baud rates or different modulation rates.

25 18. The system as in Claim 6 wherein collisions are signaled from the transmitting stations to all other stations by a duration of transmission interval that is distinguishable from the range of transmission intervals used by non-colliding transmissions.

19. The system as in Claim 18 wherein the first station simulates a collision by forcing its transmission to have a duration that falls within the predetermined interval defined for collisions.

5

20. The system as in Claim 6 wherein the communication channel is selected from the following group:

- a) baseband, utilizing frequencies down to DC;
- b) frequency division multiplexed, wherein each channel uses a distinct passband of frequencies on the physical medium;
- c) code division multiplexed, wherein each channel is associated with a set of pseudo-random code sequences; or
- d) time division multiplexed, wherein each channel occupies a specific time interval in a repeating time frame.

10  
15

21. The system as in Claim 6 wherein the physical medium of the channel is selected from the following group:

- a) twisted pair;
- b) coaxial cable;
- c) power line;
- d) optical fiber;
- e) infrared; or
- f) wireless radio frequency.

20

22. The system as in Claim 6 wherein the signal modulation used to encode packet data is selected from the following group:

- a) pulse amplitude modulation;
- b) quadrature amplitude modulation and carrierless amplitude  
5 phase modulation;
- c) multi-tone modulation; or
- d) frequency or phase modulation.

23. The system as in Claim 6 wherein a HOLDOFF signal is used to  
10 suspend the distributed fair priority queuing MAC protocol for a predetermined time interval such that the predetermined time interval can be used for access by a third station that is not executing the distributed fair priority queuing MAC protocol.

15 24. The system as in Claim 23 wherein the HOLDOFF signal and a transmission timed to end at a frame boundary is used to implement a synchronous frame interval within which the distributed fair priority queuing MAC protocol functions.

20 25. Computer data signals in a carrier wave signal transmitted over a packet-switched multiple-access network comprising:

a collision signal, the collision signal indicating a collision of a packet sent by a first station connected to the network with a packet sent by a second station connected to the network; and

25 at least two contention slots, the contention slots, which each have a different priority to provide multiple levels of priority of access, being used



for contention transmission opportunities controlled by a distributed fair priority queuing MAC (Media Access Control) protocol executed by the first station and by the second station to resolve the packet collision on a per packet per access priority basis.

5

26. The carrier wave signal as in Claim 25 further comprising:

multiple signal slots, the signal slots being used for signaling opportunities by the first station and by the second station to share information used for collision resolution by the distributed fair priority queuing MAC protocol executed by the first station and by the second station.

10

27. The carrier wave signal as in Claim 25 further comprising:

side-band signaling at the MAC layer, the side-band signaling providing information for the PHY (physical) layer; and

15

a deterministic label indicator slot for indicating that the pattern of votes is the result of a station intending to convey side-band signaling information for the PHY layer.

20

28. The carrier wave signal as in Claim 25 further comprising:

a HOLDOFF signal, the HOLDOFF signal being used to suspend the distributed fair priority queuing MAC protocol for a predetermined time interval such that the predetermined time interval can be used for access by a third station that is not executing the distributed fair priority queuing MAC protocol.

25

29. A method for a packet-switched multiple-access network, the method comprising:

transmitting packets to a shared communication channel and receiving packets from the channel; and

5       executing a distributed fair priority queuing MAC (Media Access Control) protocol, the distributed fair priority queuing MAC protocol resolving packet collisions on the channel on a per packet per access priority basis such that stations connected to the channel will have a fair and prioritized opportunity to transmit pending packets.

10

30. The method as in Claim 29 further comprising:  
detecting the presence or absence of transmissions on the channel.

31. The method as in Claim 30 further comprising:  
15       demodulating received transmissions into packet data.

32. The method as in Claim 31 further comprising:  
modulating packet data into signals for transmission.

20       33. The method as in Claim 32 further comprising:  
detecting a collision.

34. The method as in Claim 33 further comprising:

providing a plurality of contention slots, the contention slots being used for packet transmission contention opportunities that are ordered according to access priority.

5           35.    The method as in Claim 34 further comprising:  
              providing a plurality of signal slots.

              36.    The method as in Claim 35 wherein the signal slots are used for  
signaling opportunities for the first station to share information with other stations,  
10    which are connected to the channel, used for collision resolution by the distributed  
fair priority queuing MAC protocol.

              37.    The method as in Claim 35 wherein the signal slots are used for  
signaling opportunities for a first station to share information with other stations,  
15    which are connected to the channel, and the signal slots are used for collision  
resolution protocol implemented using a tree or stack-based collision resolution  
technique.

              38.    The method as in Claim 35 wherein each station contending for  
20    access in a given contention resolution cycle pseudo-randomly selects among the  
possible signal slots and signals that vote by transmitting a signal during the  
selected signal slot.

              39.    The method as in Claim 35 further comprising:

providing a first counter per priority that indicates the stack level above the current active stations with which the station is waiting to resolve its collision, and

5 incrementing and decrementing the first counter based on one or more of the following: (a) a carrier sense signal; (b) a collision detect signal; or (c) voting signals.

40. The method as in Claim 39 further comprising:

10 providing a second counter that indicates the maximum stack level per priority for any station connected to the channel, wherein the second counter is used to initialize the first counter when the first station has prepared a new packet for transmission.

15 41. The method as in Claim 35 wherein a pattern of votes used in a collision resolution cycle conveys side-band signaling information from the first station to other stations connected to the channel.

20 42. The method in Claim 35 wherein a pattern of votes used in a collision resolution cycle is used to convey a subordinate level of priority within the priority level associated with the collision resolution cycle.

25 43. The method in Claim 35 wherein a pattern of votes used in a collision resolution cycle is chosen such that the probability of any second station further colliding with the first station is reduced.

44. The method as in Claim 35 wherein a collision resolution cycle comprises two, three or four signal slots.

5 45. The method as in Claim 35 wherein collisions are signaled from transmitting stations to other stations connected to the channel by a predetermined duration of transmission interval that is distinguishable from the duration of transmission intervals used by non-colliding transmissions.

10 46. The method as in Claim 45 wherein the first station simulates a collision by forcing its transmission to have a duration that falls within the predetermined duration defined for collisions.

47. The method as in Claim 34 further comprising:  
15 transmitting a HOLDOFF signal, the HOLDOFF signal being used to suspend the distributed fair priority queuing MAC protocol for a predetermined time interval such that the predetermined time interval can be used for access by a station that is not executing the distributed fair priority queuing MAC protocol.

20 48. The method as in Claim 34 wherein the plurality of contention slots are further used for segregating transmissions at different baud rates or different modulation formats.

25 49. The method as in Claim 34 wherein data packets are transmitted over the channel and the signal slots are transmitted over a second channel.

50. The method as in Claim 41 further comprising:
- providing a deterministic label indicator slot for indicating that the pattern of votes is the result of a station intending to convey side-band signaling information.

# INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/03070

## A. CLASSIFICATION OF SUBJECT MATTER

IPC 6 H04L12/407 H04L12/413 G06F13/376

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04L G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 4 623 886 A (LIVINGSTON WILLIAM D) 18 November 1986 (1986-11-18)  abstract column 4, line 48 - column 5, line 35 column 7, line 6 - line 40 column 8, line 31 - column 9, line 31 column 10, line 54 - column 11, line 5 claims 3,5	1-7, 20-22, 25, 29-34
A	US 5 553 072 A (DAGGETT KENNETH E ET AL) 3 September 1996 (1996-09-03) abstract claim 10	21

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

\* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

19 July 1999

Date of mailing of the international search report

03/08/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Karavassilis, N

# INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/03070

## C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>EP 0 299 639 A (DAVID SYSTEMS INC)  18 January 1989 (1989-01-18)  abstract  page 3, column 3, line 28 - line 42  -----</p>	22



# INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 99/03070

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 4623886	A	18-11-1986	NONE	
US 5553072	A	03-09-1996	NONE	
EP 0299639	A	18-01-1989	NONE	